
A PAC-Bayesian Analysis of Graph Clustering and Pairwise Clustering

Yevgeny Seldin

Max Planck Institute for Biological Cybernetics
Tübingen, Germany
seldin@tuebingen.mpg.de

Abstract

We formulate weighted graph clustering as a prediction problem¹: given a subset of edge weights we analyze the ability of graph clustering to predict the remaining edge weights. This formulation enables practical and theoretical comparison of different approaches to graph clustering as well as comparison of graph clustering with other possible ways to model the graph. We adapt the PAC-Bayesian analysis of co-clustering (Seldin and Tishby, 2008; Seldin, 2009) to derive a PAC-Bayesian generalization bound for graph clustering. The bound shows that graph clustering should optimize a trade-off between empirical data fit and the mutual information that clusters preserve on the graph nodes. A similar trade-off derived from information-theoretic considerations was already shown to produce state-of-the-art results in practice (Slonim et al., 2005; Yom-Tov and Slonim, 2009). This paper supports the empirical evidence by providing a better theoretical foundation, suggesting formal generalization guarantees, and offering a more accurate way to deal with finite sample issues. We derive a bound minimization algorithm and show that it provides good results in real-life problems and that the derived PAC-Bayesian bound is reasonably tight.

1 Introduction

Graph clustering is an important tool in data analysis with wide variety of applications including social networks analysis, bioinformatics, image processing, and many more. As a result a multitude of different approaches to graph clustering were developed. Examples include graph cut methods (Shi and Malik, 2000), spectral clustering (Ng et al., 2001), information-theoretic approaches (Slonim et al., 2005), to name just a few. Comparing the different approaches is usually a painful task, mainly because the goal of each of these clustering methods is formulated in terms of the solution: most clustering methods start by defining some objective functional and then minimizing it. But for a given problem how can we choose whether to apply a graph cut method, spectral clustering, or an information-theoretic approach?

In this paper we formulate weighted graph clustering as a prediction problem². Given a subset of edge weights we analyze the ability of graph clustering to predict the remaining edge weights. The rational behind this formulation is that if a model (not necessarily cluster-based) is able to predict with high precision all edge weights of a graph given a small subset of edge weights then it is a good model of the graph. The advantage of this formulation of graph modeling is that it is independent of a specific way chosen to model the graph and can be used to compare any two solutions, either by comparison of generalization bounds or by cross-validation. The generalization bound or cross-

¹Pairwise clustering is equivalent to clustering of a weighted graph, where edge weights correspond to pairwise distances. Hence, from this point on, we restrict the discussion to graph clustering.

²Unweighted graphs can be modeled by setting the weight of present edges as 1 and absent edges as 0.

validation also address the finite-sample nature of the graph clustering problem and provide a clear criterion for model order selection. For very large datasets, where computational constraints can prevent considering all edges of a graph, as for example in (Yom-Tov and Slonim, 2009), the generalization bound can be used to resolve the trade-off between computational workload and precision of graph modeling.

The formulation and analysis of graph clustering presented here are based on the analysis of co-clustering suggested in (Seldin and Tishby, 2008; Seldin, 2009), which is reviewed briefly in section 2. In section 3 we adapt the analysis to derive PAC-Bayesian generalization bound for the graph clustering problem. The generalization bound depends on a trade-off between empirical fit of the cluster structure to the graph and the amount of mutual information that the clusters preserve on the graph nodes. This trade-off is related to the objective of a successful graph clustering algorithm Iclust (Slonim et al., 2005). We discuss this relation in section 4. In section 5 we suggest an algorithm for minimization of our bound and, finally, in section 6 we present some experiments with real-world data and analyze the tightness of the bound.

2 Review of PAC-Bayesian Analysis of Co-clustering

Co-clustering is a widely used method for analysis of data in the form of a matrix by simultaneous clustering of rows and columns of the matrix (Banerjee et al., 2007). A good illustrative example of a co-clustering problem is collaborative filtering (Herlocker et al., 2004). In collaborative filtering one is given a matrix of viewers by movies with ratings given by the viewers to the movies. The matrix is usually sparse and the task is to predict the missing entries. We assume that there is an unknown probability distribution $p(X_1, X_2, Y)$ over the triplets of viewer X_1 , movie X_2 , and rating Y . The goal is to build a discriminative predictor $q(Y|X_1, X_2)$ that given a viewer and movie pair will predict the expected rating Y . A natural form of evaluation of such predictors, no matter whether they are based on co-clustering or not, is to evaluate the expected loss $\mathbb{E}_{p(X_1, X_2, Y)} \mathbb{E}_{q(Y'|X_1, X_2)} l(Y, Y')$, where $l(Y, Y')$ is an externally provided loss function for predicting Y' instead of Y .

2.1 PAC-Bayesian Analysis of Discriminative Prediction with Co-clustering

Let $\mathcal{X}_1 \times \dots \times \mathcal{X}_d \times \mathcal{Y}$ be a $(d+1)$ -dimensional product space and assume that each \mathcal{X}_i is categorical and its cardinality $|\mathcal{X}_i|$ is fixed and known. We also assume that \mathcal{Y} is finite with cardinality $|\mathcal{Y}|$ and that the loss function $l(Y, Y')$ is bounded. In the collaborative filtering example \mathcal{X}_1 is the space of viewers, \mathcal{X}_2 is the space of movies, $d = 2$, and \mathcal{Y} is the space of ratings (e.g., on a five-star scale). The loss $l(Y, Y')$ can be, for example, an absolute loss $l(Y, Y') = |Y - Y'|$ or a quadratic loss $l(Y, Y') = (Y - Y')^2$.

We assume an existence of an unknown probability distribution $p(X_1, \dots, X_d, Y)$ over $\mathcal{X}_1 \times \dots \times \mathcal{X}_d \times \mathcal{Y}$ and that a training sample of size N is generated i.i.d. according to p . We use $\hat{p}(X_1, \dots, X_d, Y)$ to denote the empirical frequencies of $(d+1)$ -tuples $\langle X_1, \dots, X_d, Y \rangle$ in the sample. We consider the following form of discriminative predictors:

$$q(Y|X_1, \dots, X_d) = \sum_{C_1, \dots, C_d} q(Y|C_1, \dots, C_d) \prod_{i=1}^d q(C_i|X_i). \quad (1)$$

The hidden variables C_1, \dots, C_d represent a clustering of X_1, \dots, X_d . The hidden variable C_i accepts values in $\{1, \dots, |C_i|\}$, where $|C_i|$ is the number of clusters used along dimension i . The free parameters of the model (1) are the conditional probability distributions $q(C_i|X_i)$ which represent the probability of assigning X_i to cluster C_i and the conditional probability $q(Y|C_1, \dots, C_d)$ which represents the probability of assigning label Y to cell $\langle C_1, \dots, C_d \rangle$ in the cluster product space. We denote the free parameters collectively by $\mathcal{Q} = \{ \{q(C_i|X_i)\}_{i=1}^d, q(Y|C_1, \dots, C_d) \}$. We define the expected and empirical losses $L(\mathcal{Q})$ and $\hat{L}(\mathcal{Q})$ of the prediction strategy defined by \mathcal{Q} as:

$$L(\mathcal{Q}) = \mathbb{E}_{p(X_1, \dots, X_d, Y)} \mathbb{E}_{q(Y'|X_1, \dots, X_d)} l(Y, Y'), \quad (2)$$

$$\hat{L}(\mathcal{Q}) = \mathbb{E}_{\hat{p}(X_1, \dots, X_d, Y)} \mathbb{E}_{q(Y'|X_1, \dots, X_d)} l(Y, Y'), \quad (3)$$

where $q(Y|X_1, \dots, X_d)$ is defined by (1). We define the mutual information $\bar{I}(X_i; C_i)$ corresponding to the joint distribution $\bar{q}(X_i; C_i) = \frac{1}{|\mathcal{X}_i|} q(C_i|X_i)$ defined by $q(C_i|X_i)$ and a *uniform* distribution

over X_i as:

$$\bar{I}(X_i; C_i) = \frac{1}{|X_i|} \sum_{x_i \in \mathcal{X}_i} \sum_{c_i \in \mathcal{C}_i} q(c_i|x_i) \ln \frac{q(c_i|x_i)}{\bar{q}(c_i)}, \quad (4)$$

where $\bar{q}(c_i) = \frac{1}{|X_i|} \sum_{x_i} q(c_i|x_i)$ is the marginal distribution over C_i . Finally, we denote the KL-divergence between two Bernoulli distributions with biases $\hat{L}(\mathcal{Q})$ and $L(\mathcal{Q})$ by

$$kl(\hat{L}(\mathcal{Q})||L(\mathcal{Q})) = \hat{L}(\mathcal{Q}) \ln \frac{\hat{L}(\mathcal{Q})}{L(\mathcal{Q})} + (1 - \hat{L}(\mathcal{Q})) \ln \frac{1 - \hat{L}(\mathcal{Q})}{1 - L(\mathcal{Q})}. \quad (5)$$

The following generalization bound for discriminative prediction with co-clustering was proved in (Seldin, 2009).

Theorem 1. *For any probability measure $p(X_1, \dots, X_d, Y)$ over $\mathcal{X}_1 \times \dots \times \mathcal{X}_d \times \mathcal{Y}$ and for any loss function l bounded by 1, with a probability of at least $1 - \delta$ over a selection of an i.i.d. sample S of size N according to p , for all randomized classifiers $\mathcal{Q} = \{\{q(C_i|X_i)\}_{i=1}^d, q(Y|C_1, \dots, C_d)\}$:*

$$kl(\hat{L}(\mathcal{Q})||L(\mathcal{Q})) \leq \frac{\sum_{i=1}^d (|X_i| \bar{I}(X_i; C_i) + |C_i| \ln |X_i|) + \left(\prod_{i=1}^d |C_i|\right) \ln |Y| + \frac{1}{2} \ln(4N) - \ln \delta}{N}. \quad (6)$$

In practice Seldin (2009) replace (6) with a parameterized trade-off

$$\mathcal{F}(\mathcal{Q}) = \beta N \hat{L}(\mathcal{Q}) + \sum_{i=1}^d n_i \bar{I}(X_i; C_i) \quad (7)$$

and suggest an alternating projection algorithm for finding a local minimum of $\mathcal{F}(\mathcal{Q})$ (for a fixed β). Bound (6) is minimized by applying a linear search over β and substituting $\hat{L}(\mathcal{Q})$ and $\bar{I}(X_i; C_i)$ obtained from optimization of $\mathcal{F}(\mathcal{Q})$ back into (6). Alternatively, the value of β can be tuned by cross-validation. This algorithm achieved state-of-the-art performance on the MovieLens collaborative filtering dataset. Below we adapt this analysis and algorithm to the graph clustering problem.

3 Formulation and Analysis of Graph Clustering

3.1 Graph Clustering as a Prediction Problem

Assume that \mathcal{X} is a space of $|X|$ nodes and denote by w_{ij} the weight of an edge connecting nodes i and j .³ We assume that the weights w_{ij} are generated according to an unknown probability distribution $p(W|X_1, X_2)$, where $X_1, X_2 \in \mathcal{X}$ are the edge endpoints. We further assume that we know the space of nodes \mathcal{X} and are given a sample of size N of edge weights, generated according to $p(X_1, X_2, W)$. The goal is to build a regression function $q(W|X_1, X_2)$ that will minimize the expected prediction error of the edge weights $\mathbb{E}_{p(X_1, X_2, W)} \mathbb{E}_{q(W'|X_1, X_2)} l(W, W')$ for some externally given loss function $l(W, W')$. Note that this formulation does not assume any specific form of $q(W|X_1, X_2)$ and enables comparison of all possible approaches to this problem.

3.2 PAC-Bayesian Analysis of Graph Clustering

In this work we analyze the generalization abilities of $q(W|X_1, X_2)$ based on clustering:

$$q(W|X_1, X_2) = \sum_{C_1, C_2} q(W|C_1, C_2) q(C_1|X_1) q(C_2|X_2). \quad (8)$$

One can immediately see the relation between (8) and (1). The only difference is that in (8) the nodes X_1, X_2 belong to the same space of nodes \mathcal{X} and the conditional distribution $q(C|X)$ is shared for the mapping of endpoints of an edge. Let $\hat{p}(X_1, X_2, W)$ be the empirical distribution over edge

³ All the results can be straightforwardly extended to hyper-graphs.

weights. The empirical loss of a prediction strategy $\mathcal{Q} = \{q(C|X), q(W|C_1, C_2)\}$ corresponding to (8) can then be written as:

$$\hat{L}(\mathcal{Q}) = \mathbb{E}_{\hat{p}(X_1, X_2, W)} \mathbb{E}_{q(W'|X_1, X_2)} l(W, W'). \quad (9)$$

The following generalization bound for graph clustering can be proved by a minor adaptation of the proof of theorem 1.

Theorem 2. *For any probability measure $p(X_1, X_2, W)$ over the space of nodes and edge weights $\mathcal{X} \times \mathcal{X} \times \mathcal{W}$ and for any loss function l bounded by 1, with a probability of at least $1 - \delta$ over a selection of an i.i.d. sample S of size N according to p , for all graph clustering models defined by $\mathcal{Q} = \{q(C|X), q(W|C_1, C_2)\}$:*

$$kl(\hat{L}(\mathcal{Q}) \| L(\mathcal{Q})) \leq \frac{|X|\bar{I}(X; C) + |C| \ln |X| + |C|^2 \ln |W| + \frac{1}{2} \ln(4N) - \ln \delta}{N}, \quad (10)$$

where $|C|$ is the number of node clusters and $|W|$ is the number of distinct edge weights.

The limitation of working with a fixed set of allowed edge weights is resolved by weight quantization in section 5.1.

Although there is no analytical expression for the inverse KL-divergence, given (10) we can easily bound $L(\mathcal{Q})$ numerically:

$$\begin{aligned} L(\mathcal{Q}) &\leq kl^{-1} \left(\hat{L}(\mathcal{Q}), \frac{|X|\bar{I}(X; C) + |C| \ln |X| + |C|^2 \ln |W| + \frac{1}{2} \ln(4N) - \ln \delta}{N} \right) \\ &= \max \left\{ z : kl(\hat{L}(\mathcal{Q}) \| z) \leq \frac{|X|\bar{I}(X; C) + |C| \ln |X| + |C|^2 \ln |W| + \frac{1}{2} \ln(4N) - \ln \delta}{N} \right\}. \end{aligned} \quad (11)$$

Similar to the approach applied by Seldin (2009) in co-clustering, in practice we can replace (10) with a parameterized trade-off:

$$\mathcal{G}(\mathcal{Q}) = \beta N \hat{L}(\mathcal{Q}) + |X| \bar{I}(X; C) \quad (12)$$

and tune β either by substituting $\hat{L}(\mathcal{Q})$ and $\bar{I}(X; C)$ resulting from a solution of (12) back into (11) or via cross-validation. In section 5 we suggest an algorithm for minimization of (12).

4 Related Work

The regularization of pairwise clustering by mutual information $\bar{I}(X; C)$ was already applied in practice by Slonim et al. (2005). In their work they maximized a parameterized trade-off $\langle s \rangle - T\bar{I}(X; C)$, where $\langle s \rangle = \sum_c \bar{q}(c) \sum_{x_1, x_2} q(x_1|c)q(x_2|c)w_{x_1 x_2}$ measured average pairwise similarities within a cluster⁴. Their algorithm demonstrated superior results in cluster coherence compared to 18 other clustering methods. The regularization by mutual information was motivated by information-theoretic considerations inspired by the rate distortion theory (Cover and Thomas, 1991). Namely, the authors drew a parallel between $\langle s \rangle$ and distortion and $\bar{I}(X; C)$ and compression rate of a clustering algorithm. Further, Yom-Tov and Slonim (2009) showed that the algorithm can be run in parallel mode, where each parallel worker operates with a subset of pairwise relations at each iteration rather than all of them. Such mode of operation was motivated by inability to consider all pairwise relations in very large datasets due to computational constraints. Yom-Tov and Slonim (2009) reported only minor empirical degradation in clustering quality, but no formal analysis and guarantees were suggested.

In light of this prior work the main contribution of our paper is not as much the introduction of the trade-off $\mathcal{G}(\mathcal{Q})$ in equation (12), but rather the formulation of graph clustering as a prediction problem and the analysis of the finite sample aspect of this problem. The experiments that follow focus on the analysis of tightness of the bound derived in section 3.

⁴The loss $L(\mathcal{Q})$ is slightly more general than $\langle s \rangle$ since it also considers edges between the clusters.

5 An Algorithm for Graph Clustering

In this section we derive an algorithm for minimization of the trade-off $\mathcal{G}(\mathcal{Q})$. Unlike the co-clustering trade-off $\mathcal{F}(\mathcal{Q})$ in equation (7), which is convex in $q(C_1|X_1)$ and $q(C_2|X_2)$ and thus can be minimized by alternating projections, the trade-off $\mathcal{G}(\mathcal{Q})$ is not convex in $q(C|X)$. Nevertheless, we found in our experiments that alternating projections still provide good outcome in practice. Alternatively, one can apply sequential minimization techniques, as done by Yom-Tov and Slonim (2009). The alternating projections are much faster though and for that reason were chosen for the experiments.

The alternating projections are derived similar to alternating projection minimization in the rate distortion theory (Cover and Thomas, 1991), namely by writing the Lagrangian corresponding to $\mathcal{G}(\mathcal{Q})$, deriving it with respect to the free parameters and equating the derivative to zero. This procedure provides a set of self-consistent equations, which are exactly the same as those for alternating projection of $\mathcal{F}(\mathcal{Q})$, hence we write the result in the Algorithm 1 box and refer the reader to (Seldin, 2009) for derivation details. The only difference in our case is in the form of the derivative $\frac{\partial \hat{L}(\mathcal{Q})}{\partial q(c|x)}$, which we derive next.

Algorithm 1 One iteration of an alternating projection of $\mathcal{G}(\mathcal{Q}) = \beta N \hat{L}(\mathcal{Q}) + |X| \bar{I}(X; C)$.

Input: $\hat{p}(x_1, x_2, w)$, $q_t(C|X)$, $g_t(c_1, c_2)$, N , $|X|$, $|C|$, $l(w, w')$, β .
 $\bar{q}_t(c) \leftarrow \frac{1}{|X|} \sum_x q_t(c|x)$
 $q_{t+1}(c|x) \leftarrow \bar{q}_t(c) e^{-\beta N \frac{\partial \hat{L}(\mathcal{Q}_t)}{\partial q(c|x)}}$
 $Z_{t+1}(x) \leftarrow \sum_c q_{t+1}(c|x)$
 $q_{t+1}(c|x) \leftarrow \frac{q_{t+1}(c|x)}{Z_{t+1}(x)}$
 $g_{t+1}(c_1, c_2) \leftarrow \arg \min_{w'} \sum_w l(w, w') \sum_{x_1, x_2} q_{t+1}(c_1|x_1) \hat{p}(x_1, x_2, w) q_{t+1}(c_2|x_2)$
return $q_{t+1}(C|X)$, $g_t(C_1, C_2)$.

For notational convenience we reformulate the problem in matrix notation. For simplicity we assume that the edge weights w are sampled without repetition. This assumption usually holds in practice and it also does not affect the tightness of the analysis since the convergence rate of sampling without repetition is lower bounded by the convergence rate of sampling with repetition (Derbeko et al., 2004). With this assumption we can represent the training data by the Hadamard (also known as Schur) entrywise matrix product $S \circ M$ (denoted by $S .* W$ in Matlab), where $S_{ij} = 1$ if the edge from node i to node j was observed in the sample and $S_{ij} = 0$ otherwise, and $W_{ij} = w_{ij}$. In order to obtain the derivative $\frac{\partial \hat{L}(\mathcal{Q})}{\partial q(c|x)}$ we have to assume a specific form of $l(w, w')$. We choose quadratic loss $l(w, w') = (w - w')^2$. The maximum likelihood reconstruction (the one that minimizes $\hat{L}(\mathcal{Q})$) for the quadratic loss is a delta distribution $q(w|c_1, c_2) = \delta(w, g(c_1, c_2))$, where $g(c_1, c_2) = \arg \min_{w'} \sum_w l(w, w') \sum_{x_1, x_2} q(c_1|x_1) \hat{p}(x_1, x_2, w) q(c_2|x_2) = \sum_{x_1, x_2, w} q(c_1|x_1) w \hat{p}(x_1, x_2, w) q(c_2|x_2)$. This enables us to write the prediction model (8) and the loss $\hat{L}(\mathcal{Q})$ in a matrix form. Let Q be the matrix of $q(c|x)$ with rows indexed by cluster variables and columns indexed by node variables and G be the matrix of weights predicted in the cluster product space. We denote the elements of G by $g(c_1, c_2)$. The prediction model (8) can then be written as

$$g(x_1, x_2) = \sum_{c_1, c_2} q(c_1|x_1) g(c_1, c_2) q(c_2|x_2) \quad (13)$$

and the corresponding reconstruction matrix is $Q^T G Q$. Note that $g(x_1, x_2)$ is a function of x_1, x_2 , which corresponds to a probability distribution $q(w|x_1, x_2)$, which is a delta function. The loss can then be written as:

$$\hat{L}(\mathcal{Q}) = \frac{1}{N} \|S \circ (M - Q^T G Q)\|_2^2, \quad (14)$$

where $\|\cdot\|_2^2$ is the squared Frobenius norm of a matrix. The maximum likelihood G is given by $G = Q(S \circ M)Q^T/N$ and the derivative $\frac{\partial \hat{L}(\mathcal{Q})}{\partial q(c|x)} = 4G^T Q(S \circ (Q^T G Q - M))/N$.

Equation (14) provides an easy way to see why $\hat{L}(\mathcal{Q})$ and hence $\mathcal{G}(\mathcal{Q})$ are not convex in Q - since Q appears in forth power. Therefore, repeated iteration of alternative projections in Algorithm 1 is not

guaranteed to converge (and indeed it does not). However, we found that empirically even a single iteration of Algorithm 1 achieves remarkably good results and due to simplicity of the algorithm it is easy to try multiple random initializations and obtain results comparable to those obtained by sequential optimization within much shorter time. This was the strategy followed in this paper. For large number of clusters we found it useful to anneal β from a lower value $\beta' = 1/N$ up to the desired value in two-fold increments. At each value of β we iterated alternating projections for 5 times and then added a small random noise to $q(c|x)$ before increasing β by a factor of 2 until reaching the desired value.

5.1 Correction for Edge Weight Quantization

We note that the alternating projections algorithm derived above operates with continuous weights w , whereas the analysis in theorem 2 allows only a finite set of edge weights. If the edge weights are uniformly quantized at intervals Δ , then $|W| = \frac{1}{\Delta}$ (assume that the quantization starts at $\frac{1}{2}\Delta$ and ends at $1 - \frac{1}{2}\Delta$). By rounding the continuous edge weights obtained by the alternating projections toward the closest quantization both the empirical and the expected loss are increased by at most $\Delta + \frac{1}{4}\Delta^2$. This is because quantization can shift the prediction by at most $\frac{1}{2}\Delta$ and then $l(w, w' + \frac{1}{2}\Delta) = (w - w' - \frac{1}{2}\Delta)^2 = (w - w')^2 - (w - w')\Delta + \frac{1}{4}\Delta^2 \leq l(w, w') + \Delta + \frac{1}{4}\Delta^2$, where the last inequality follows from the assumption that the loss $l(w, w')$ is bounded by 1. Hence, for the continuous weights we have

$$L(\mathcal{Q}) \leq kl^{-1} \left(\hat{L}(\mathcal{Q}) + \Delta + \frac{\Delta^2}{4}, \frac{|X|\bar{I}(X; C) + |C| \ln |X| - |C|^2 \ln \Delta + \frac{1}{2} \ln \frac{4N}{\delta^2}}{N} \right) + \Delta + \frac{\Delta^2}{4}. \quad (15)$$

As a rule of thumb we have taken $\Delta = 5|C|^2/N$, so that the contribution of Δ to the two operands of the inverse KL-divergence is approximately equivalent. In general this correction for quantization had no significant influence on the bound.

6 Applications

We evaluate the bound derived in section 3 and the algorithm for its minimization from section 5 on two real-life datasets used previously in (Yom-Tov and Slonim, 2009). The first dataset named “king” was taken from (Gummadi et al., 2002). The graph represents a set of 1,740 DNS servers and the edge weights correspond to similarities between the servers. The similarities are negative exponents of the latencies between the servers scaled by dividing by the median value of all latencies in the data. The second dataset contained the graph of all known pairwise interactions among 5,202 Yeast proteins, downloaded on February 15, 2008 from the BioGRID web site⁵. The edge weights were set to be 1 between interacting proteins and 0 otherwise.

King Dataset Experiments

In the first experiment we split the king dataset into five random train, cross-validation, and test subsets. The train set size is 103,866 edge weights, the cross-validation set size is 25,967 edges and the test set consists of the remaining 1,383,097 edges. The size of the train set is only 3.4% of all edges or if compared to the size of the node space the number of observed edges is $8|X| \ln |X|$. This level of sparsity is even slightly lower than the 5.3% fraction of edges considered in each iteration of the parallel Iclust algorithm in (Yom-Tov and Slonim, 2009) (the total number of edges considered in all iterations of parallel Iclust was generally larger). We cluster the graph into 41 clusters, which is the same number used by Yom-Tov and Slonim (2009) and compare the test loss and the value of bound (15) as a function of β . I.e., for each value of β we minimize $\mathcal{G}(\mathcal{Q})$ using the alternating projections algorithm and substitute the resulting $\hat{L}(\mathcal{Q})$ and $\bar{I}(X; C)$ into (15) to compute the bound. The result is shown in Figure 1.a. The bound is not perfectly tight, mainly due to the large $|C|^2 \ln |W|$ term in this case. Nevertheless, the bound is meaningful and the cross-validation loss almost coincides with the test loss.

In the second experiment we consider all edges and cluster the dataset into $|C| = 1, 2, \dots, 15$ clusters. (Due to symmetry every edge in this dataset appears twice, once from node i to j and another time

⁵<http://www.thebiogrid.org/downloads.php>

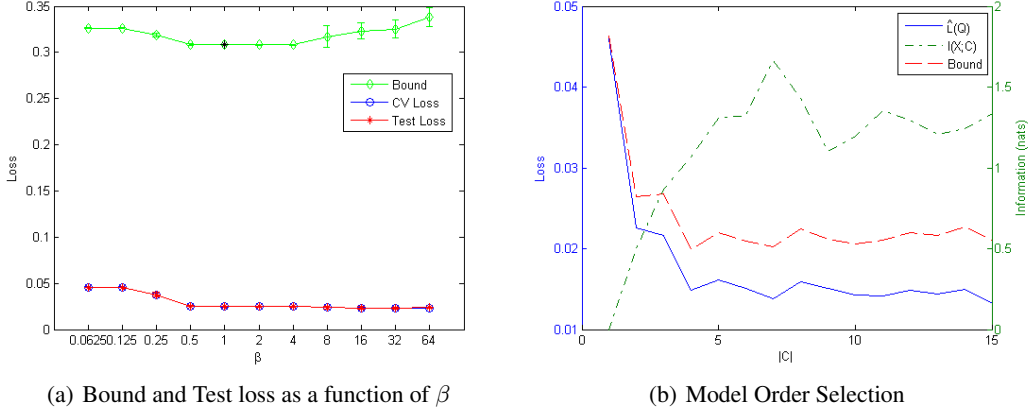


Figure 1: **King dataset experiments.** (a) Bound (15), cross-validation loss, and test loss as a function of β . Error bars indicate one standard deviation. The minimum of the bound is indicated by the black “*”. Cross-validation follows the test loss so closely, that the curves coincide. (b) Train loss $\hat{L}(Q)$, information $\bar{I}(X; C)$, and bound (15) as a function of $|C|$.

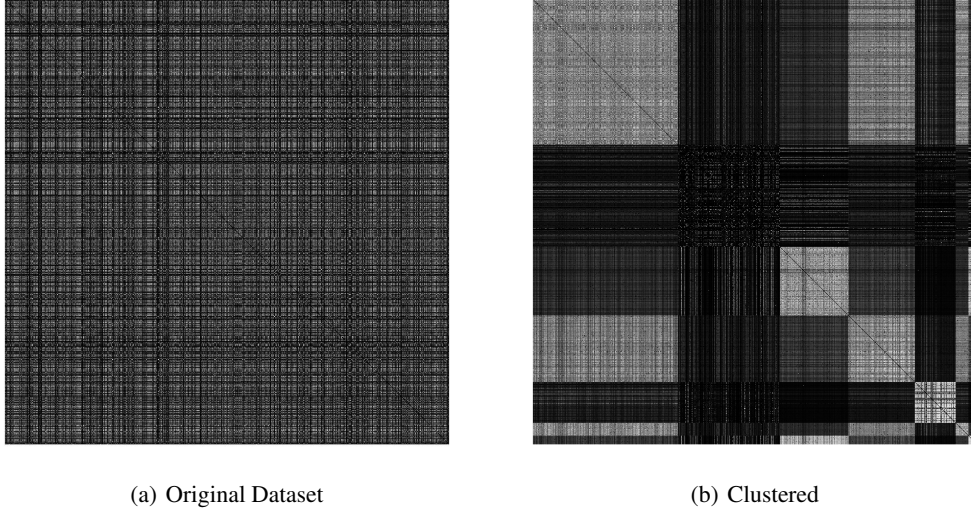


Figure 2: **Illustration of the king dataset.** (a) Original dataset. (b) Clustering into 7 clusters.

from node j to i , but in our analysis we consider only one copy of each edge.) The value of β in the optimization trade-off $\mathcal{G}(Q)$ was set to 1. In general by search for the optimal β the results could be improved slightly, although as we can see from the previous experiment not considerably, so we omitted the search over β in this experiment. The results are shown in Figure 1.b. First, we see that modeling this dataset by clustering is provably beneficial: the expected loss in predicting the weights of missing edges (would there be any) drops from 0.046 when predicting the weight with the global average to 0.02 when using four clusters and remains roughly at this level when the number of clusters is further increased. To the best of our knowledge, this is the first time when the benefit of clustering is formally proven and measured without any assumptions on the distribution that generated the edge weights (except that they were generated independently from that distribution). In this experiment there is no test set, but we can see that the bound follows the train loss pretty tightly. The mutual information preserved by the clusters on the node variables saturates at about 1.2-1.5 nats, which corresponds to effective complexity of about four clusters. Clustering of the dataset into seven clusters is illustrated in Figure 2.

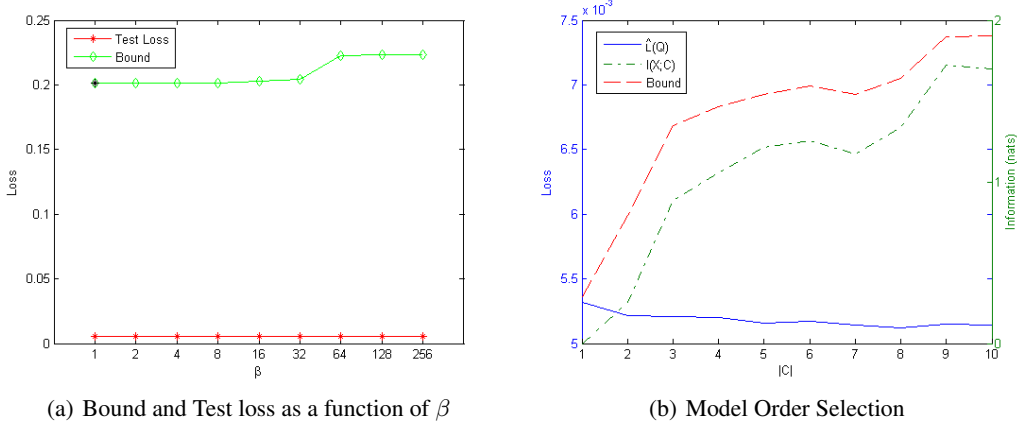


Figure 3: **Yeast dataset experiments.** (a) Bound (15) and test loss as a function of β . The minimum of the bound is indicated by the black “*”. (b) Train loss $\hat{L}(\mathcal{Q})$, information $\bar{I}(X; C)$, and bound (15) as a function of $|C|$. Note that the bound scale is on the left hand side.

Yeast Dataset Experiments

In our first experiment we apply five random splits of the dataset into 445,125 training and 13,082,676 test edges. Training edges constitute only 3.3% of all the edges or $10|X|\ln|X|$ if compared to the number of graph nodes. As previously, the train set sparsity is slightly lower than the 5.3% sparsity considered in each iteration of the parallel Iclust algorithm in (Yom-Tov and Slonim, 2009). We cluster the graph into 71 clusters, which is the same number as used by Yom-Tov and Slonim (2009). The comparison of test loss with the value of the bound is presented in Figure 3.a. The bound is not perfectly tight, mainly due to the $|C|^2 \ln|W|$ term, but is still meaningful.

In our second experiment we consider all edges and cluster the graph into $|C| = 1, \dots, 10$ clusters. (Symmetric edges from i to j and from j to i were considered only once.) The value of β was set to 256. The results are shown in Figure 3.b. As with the king dataset experiment, the results could be slightly improved by optimizing β , however even for the large value of β chosen the empirical loss $\hat{L}(\mathcal{Q})$ exhibits very minor decrease as the number of clusters grows, hence the results would not change considerably by tuning β . Due to lower number of clusters and larger training set the bound is much tighter here than in the first yeast experiment (note that the bound y scale is on the left hand side of the graph). Unlike in the king experiment the bound tells that clustering does not help in modeling this dataset.

7 Discussion

We have formulated graph clustering as a prediction problem. This formulation enables direct comparison of graph clustering with any other approach to modeling the graph. By applying PAC-Bayesian analysis we have shown that graph clustering should optimize a trade-off between empirical fit of the observed graph and the mutual information that clusters preserve on the graph nodes. Prior work of Slonim et al. (2005) and Yom-Tov and Slonim (2009) underscores practical benefits of such regularization. Our formulation suggests a better founded and accurate way of dealing with the finite sample nature of the graph clustering problem and tuning the trade-off between model fit and model complexity. It also suggests formal guarantees on the approximation quality. In particular such guarantees can be used for optimization of a trade-off between approximation precision and computational workload in processing of very large datasets. Our experiments show that the bound is reasonably tight for practical purposes.

References

Arindam Banerjee, Inderjit Dhillon, Joydeep Ghosh, Srujana Merugu, and Dhamendra Modha. A generalized maximum entropy approach to Bregman co-clustering and matrix approximation.

- Journal of Machine Learning Research*, 8, 2007.
- Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. John Wiley & Sons, 1991.
- Philip Derbeko, Ran El-Yaniv, and Ron Meir. Explicit learning curves for transduction and application to clustering and compression algorithms. *Journal of Artificial Intelligence Research*, 22, 2004.
- Krishna P. Gummadi, Stefan Saroiu, and Steven D. Gribble. King: estimating latency between arbitrary internet end hosts. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurement (IMW-2002)*, 2002.
- Jonathan Herlocker, Joseph Konstan, Loren Terveen, and John Riedl. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems*, 22(1), 2004.
- Andrew Y. Ng, Michael I. Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. In *Advances in Neural Information Processing Systems (NIPS)*, 2001.
- Yevgeny Seldin. *A PAC-Bayesian Approach to Structure Learning*. PhD thesis, The Hebrew University of Jerusalem, 2009.
- Yevgeny Seldin and Naftali Tishby. Multi-classification by categorical features via clustering. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2008.
- Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8), 2000.
- Noam Slonim, Gurinder Singh Atwal, Gasper Tracik, and William Bialek. Information-based clustering. *Proceedings of the National Academy of Science*, 102(51), 2005.
- Elad Yom-Tov and Noam Slonim. Parallel pairwise clustering. In *SIAM International Conference on Data Mining (SDM)*, 2009.